# Integral Transformation with Low-Order Scaling for Large Local Second-Order Møller–Plesset Calculations

**GUNTRAM RAUHUT,[1,*] PETER PULAY,[1] HANS-JOACHIM WERNER[2]**

[1] *Department of Chemistry and Biochemistry, University of Arkansas, Fayetteville, Arkansas 72701*
[2] *Institut für Theoretische Chemie, Universität Stuttgart, Pfaffenwaldring 55, 70569 Stuttgart, Germany*

*Received 29 October 1997; accepted 5 March 1998*

**ABSTRACT:** An algorithm is presented for the four-index transformation of electron repulsion integrals to a localized molecular orbital (MO) basis. Unlike in most programs, the first two indices are transformed in a single step. This and the localization of the orbitals allows the efficient neglect of small contributions at several points in the algorithm, leading to significant time savings. Thresholds are applied to the following quantities: distant orbital pairs, the virtual space before and after the orthogonalizing projection to the occupied space, and small contributions in the transformation. A series of calculations on medium-sized molecules has been used to determine appropriate thresholds that keep the truncation errors small (below 0.01% of the correlation energy in most cases). Benchmarks for local second-order Møller–Plesset perturbation theory (MP2; i.e., MP2 with a localized MO basis in the occupied subspace) are presented for several large molecules with no symmetry, up to 975 contracted basis functions, and 60 atoms. These are among the largest MP2 calculations performed on a single processor. The computational time (with constant basis set) scales with a

*Correspondence to*: G. Rauhut
[*] *Present address*: Institut für Theoretische Chemie, Universität Stuttgart, Pfaffenwaldring 55, 70569 Stuttgart, Germany
Contract/grant sponsors: Deutsche Forschungsgemeinschaft; Fonds der Chemischen Industrie
Contract/grant sponsor: Air Force Office for Scientific Research; contract/grant number: F49620-94-1
Contract/grant sponsor: National Science Foundation; contract/grant numbers: CHE-9319929, CHE-9707202
Contract/grant sponsor: European Union TMR; contract/grant number: FMRX-CT96-088



CCC 0192-8651 / 98 / 111241-14

somewhat lower than cubic power of the molecular size, and the memory demand is moderate even for large molecules, making calculations that require a supercomputer for the traditional MP2 feasible on workstations. © 1998 John Wiley & Sons, Inc. J Comput Chem 19: 1241–1254, 1998

**Keywords:** integral transformation; low-order scaling; second-order Møller–Plesset calculations

## Introduction

For low cost post-Hartree–Fock methods (e.g., second-order Møller–Plesset perturbation theory, MP2), the transformation of electron repulsion integrals (ERIs) from an atomic orbital (AO) to molecular orbital (MO) basis is the most demanding step, in regard to both computing time and disk space. Thus, there have been numerous suggestions to improve this transformation.[1–16] Parallel implementations[12–16] have revived the topic quite recently. For MP2, only exchange integrals of the form $(ai|jb)$ are necessary ($i$, $j$ denote occupied orbitals and $a$, $b$ virtuals), and in the following we will restrict ourselves to this type. Almost all proposed algorithms have in common that the four-index transformation is split into four quarter transformations[11]:

$$(ai|jb) = \sum_{\mu} C_{\mu a}\left( \sum_{\nu} C_{\nu b}\left( \sum_{\rho} C_{\rho i}\left( \sum_{\sigma} C_{\sigma j}(\mu\rho \mid \sigma\nu)\right)\right)\right). \quad (1)$$

This partitioning reduces the formal scaling from $O(n^2V^2N^4) \approx O(n^2N^6)$ to $O(nN^4)$. In the above, **C** is the matrix of MO coefficients; $\mu$, $\nu$, $\rho$, and $\sigma$ denote AO labels; $N$ is the number of contracted basis functions; $V \approx N$ is the number of virtual orbitals; and $n$ is the number of occupied orbitals. As $n \ll N$ (ratios $N/n \approx 5$–10 are common), most of the computational effort is required in the first quarter transformation, which requires $nO(I)$ operations, where $O(I)$ is proportional to the number of nonneglected AO integrals [in practice $O(I)$ is usually between $N^2$ and $N^3$.] The main bottleneck of the traditional algorithm is the memory requirement. The fastest algorithm, which fully exploits the eight-fold permutational symmetry of the integrals,[5] still requires the storage of all half-transformed integrals in high-speed memory ($n^2N^2/2$ memory locations) and needs in addition $nSN^2$ memory locations for intermediate storage of quarter transformed integrals. ($S$ is the maximum shell size and is independent of the molecular size.) The storage for the quarter transformed integrals can be reduced to $nS^2N$ at the cost of two integral evaluations, but because all transformed integrals must be held in memory, the total memory still scales with the fourth power of the molecular size. Alternatively, one can reduce the total memory requirement to $n^2S^2$, at the cost of four integral evaluations and a resorting of the half-transformed integrals, which requires $n^2N^2/2$ words of disk space.[7,17] This approach is limited by disk space and is probably the best conventional method for calculations up to about 1000 basis functions and 100 occupied orbitals in $C_1$ symmetry. However, it is evident that for such large cases alternative approaches have to be investigated.

In this article we explore the simultaneous transformation of two indices in the first half-transformation,

$$(\overline{\mathbf{K}}^{ij})_{\mu\nu} = (\mu i|j\nu) = \sum_{\rho\sigma} C_{\rho i}C_{\sigma j}(\mu\rho|\sigma\nu). \quad (2)$$

This method was analyzed earlier by Taylor,[6] who notes that the optimum memory for a two-index transformation is $n^2N^2/2$ but the minimum memory demand is only $N^2$, although at the considerable cost of $n^2/2$ integral evaluations. The computational work in the simultaneous two-index transformation scales formally more steeply [$n^2O(I)$] than in a series of consecutive one-index transformations. However, although few of the orbital coefficients are negligible, many of them are small and a significant fraction of their products is then very small. Precomputing upper limits for these products allows an efficient prescreening for both the integral evaluation and the transformation. As will be shown, this offsets the poorer formal scaling for large systems.

This method is particularly useful if localized MOs are used. First, the correlation of electrons in distant localized orbitals (''distant pairs'') is a weak dispersionlike interaction, and its leading term de-

creases with the inverse sixth power of the distance between the orbital centers. Depending on this distance, very distant pairs can be neglected, calculated approximately,[18] or even modeled empirically.[19] This asymptotically reduces the number of orbital pairs $(ij)$ in eq. (2) from $O(n^2)$ to $O(n)$, and the storage from $n^2N^2/2$ to $f \cdot nN^2$, where $f \cdot n$ is the number of correlated orbital pairs (simply pairs in the following). For large molecules $f$ is constant, and thus $f \cdot n$ should grow linearly with the molecular size. Second, for sufficiently large molecules the number of AOs $(\mu, \nu)$ [see eq. (2)] needed for an appropriate description of a given orbital pair $(ij)$ becomes constant and independent of the molecular size as well. This reduces both the computational work and the storage requirement. Denoting the average size of the local AO space per pair by $L$, the memory requirement is reduced to $f \cdot nL^2$; asymptotically it scales only linearly with the molecular size, even for a single integral pass computation.

The method described in the present work is based on the local correlation approach introduced by Saebø and Pulay.[20–22] It was originally conceived for higher order correlation methods, and thus little attention was paid to the integral transformation. By contrast, the present article focuses on MP2 calculations and the associated integral transformation when applied to large systems by using an integral-direct approach.[23] The local correlation method has been extended to the full coupled cluster singles and doubles level[24] and coupled with the pseudospectral formalism.[25,26] Analytical local MP2 gradients have been developed recently.[27] Other, somewhat different local correlation methods are also being pursued.[28,29] The *weak orthogonality concept* of Szalewicz et al.[30] could be used within the framework of local MP2 theory to avoid the explicit projection of the AO basis. We have not yet explored this possibility. An alternative approach to the application of the local correlation concept to MP2 theory is the use of Laplace transform orbitals as presented by Häser and Almlöf.[31] This method has recently been compared with the theory outlined above.[32,33] Approximating electron repulsion integrals can also be used to speed up MP2 calculations. This resolution of identity (RI) approach was explored by various researchers.[34–37] Although this approach retains the formal fifth power scaling of the traditional MP2 method, it can lead to significant speedups (up to an order of magnitude), depending on the completeness of the auxiliary basis set used to approximate the charge densities.

## Local MP2 Theory

In the local MP2 method,[20] the Fock matrix is nondiagonal; thus, the pair correlation coefficients (amplitudes) $T_{pq}^{ij}$ have to be determined from a coupled set of linear equations. In a compact matrix formalism,[38] using generator state spin adaptation for closed shells,[39] these equations are

$$\mathbf{R}^{ij} = \mathbf{K}^{ij} + \mathbf{F}\mathbf{T}^{ij}\mathbf{S} + \mathbf{S}\mathbf{T}^{ij}\mathbf{F} - \sum_k \mathbf{S}\left[f_{ik}\mathbf{T}^{kj} + f_{kj}\mathbf{T}^{ik}\right]\mathbf{S} = \mathbf{0}. \quad (3)$$

Here $i, j, k$ are internal (occupied) MO labels, $\mathbf{R}^{ij}$ are residuum matrices, and $\mathbf{F}$ and $\mathbf{S}$ are the Fock and overlap matrices, respectively. $f_{ik}$ is the Fock matrix element between orbitals $i$ and $k$, and $\mathbf{K}^{ij}$ is the internal exchange matrix of the orbital pair $ij$.

In eq. (3) all matrices are defined in a basis of projected AOs $|p\rangle$, which are orthogonalized against the occupied space but nonorthogonal among themselves. In terms of the basis functions (AOs) $|\mu\rangle$ the projected orbitals are defined as

$$|p\rangle = \sum_\mu P_{\mu p}|\mu\rangle. \quad (4)$$

The projection matrix given by

$$\mathbf{P} = \mathbf{1} - \tfrac{1}{2}\overline{\mathbf{D}}\overline{\mathbf{S}}, \quad (5)$$

where $\overline{\mathbf{D}}$ is the closed-shell density matrix and $\overline{\mathbf{S}}$ the overlap matrix in the (unprojected) AO basis. The exchange matrices $\mathbf{K}^{ij}$ in the projected basis are defined as

$$K_{pq}^{ij} = (pi|jq) = \sum_{\mu\nu} P_{\mu p}\overline{K}_{\mu\nu}^{ij}P_{\nu q}, \quad (6)$$

where $\overline{K}_{\mu\nu}^{ij} = (\mu i|j\nu)$ are the half-transformed integrals defined in eq. (2). The transformation of $\overline{\mathbf{F}}$ and $\overline{\mathbf{S}}$ into the projected basis is analogous. In the following, eq. (6) is denoted as the second half-transformation.

Due to the coupling terms introduced by the nondiagonal Fock matrix elements $F_{pq}$ and $f_{ik}$, the set of eq. (3) has to be solved iteratively. To achieve fast convergence, it is necessary to transform the residuum matrices $\mathbf{R}^{ij}$ to pair-adapted intermedi-

ate pseudocanonical MO basis sets, perform the updates of the amplitude matrices $\mathbf{T}^{ij}$ in this form, and transform them back to the original projected nonorthogonal AO basis.[21] This procedure has been recently described in detail for the local coupled cluster method.[24] Because the iteration is performed in the small projected basis and only the amplitudes have to reside readily available in memory, there are essentially no memory limitations in this step. In our implementation of the iteration we used a *dynamic* update procedure (i.e., each amplitude matrix is updated immediately after the evaluation of its corresponding residue). This has the advantage that the residuum matrices do not have to be stored, and it also converges significantly faster than *static* updating in which *all* amplitudes are updated at once. DIIS convergence acceleration[40] cannot be applied with dynamic updating, but sorting the pairs according to their coupling to other pairs accelerates convergence efficiently in a very simple manner. In most cases 5–6 iterations are sufficient to achieve an accuracy of $5 \times 10^{-8} E_h$, which is the default convergence threshold applied in all calculations and timings presented below.

Using the residual matrices $\mathbf{R}^{ij}$ and amplitudes $\mathbf{T}^{ij}$ one can compute the Hylleraas functional

$$E_{\text{corr}} = \sum_{ij} \sum_{pq} \left( K_{pq}^{ij} + R_{pq}^{ij} \right) \tilde{T}_{pq}^{ij}, \qquad (7)$$

where the contravariant coefficient matrices $\tilde{\mathbf{T}}^{ij}$ are defined as

$$\tilde{\mathbf{T}}^{ij} = 2\mathbf{T}^{ij} - \mathbf{T}^{ji} \quad \text{with} \quad \mathbf{T}^{ji} = [\mathbf{T}^{ij}]^T. \qquad (8)$$

The Hylleraas functional gives for any choice of the amplitudes an upper bound to the MP2 energy. For the *converged* amplitudes $\mathbf{T}^{ij}$ the residual matrices $\mathbf{R}^{ij}$ vanish and the MP2 correlation energy is obtained simply as

$$E_{\text{corr}}^2 = \sum_{ij} \sum_{pq} K_{pq}^{ij} \tilde{T}_{pq}^{ij} = \sum_{i \geq j} \frac{2}{1 + \delta_{ij}} \sum_{pq} K_{pq}^{ij} \tilde{T}_{pq}^{ij} = \sum_{i \geq j} \epsilon_{ij}. \qquad (9)$$

Without further approximations (i.e., no restriction of the virtual space, etc.), these formulas are exactly equivalent to canonical MP2 theory; and the total correlation energies obtained with the local and canonical treatments agree within numerical accuracy. The individual pair energies $\epsilon_{ij}$ are different, however.

In the local correlation approach two distinct approximations are introduced. First, for a pair $(ij)$ substitutions are limited to projected orbitals $p, q$ that lie in the vicinity of the two localized orbitals $i$ and $j$. This leads to very sparse amplitude matrices, because only those elements $T_{pq}^{ij}$ are nonzero for which $p, q$ belong to the *pair domain* $[ij]$. Because the same restriction applies to the residual matrices $R_{pq}^{ij}$ in eq. (3), only elements $K_{pq}^{ij}$ with $p, q \in [ij]$ are needed. This fact can be exploited to reduce the memory requirements and the computational effort in the integral transformation. The restriction of the configuration space to projected functions in direct neighborhood of the correlated orbitals $i$ and $j$ may thus be considered as a physically well-defined configuration selection scheme. This truncation of the correlation space causes a small loss (1–2%) in the correlation energy.[22] The fraction of correlation energy loss decreases with increasing basis set quality and can be at least partially interpreted as arising from removing the intramolecular basis set superposition error.[41] Second, distant pairs can be neglected; the exchange matrices $K_{pq}^{ij}$ need to be computed only for a subset of orbital pairs $(ij)$. As will be discussed in the next section, this leads to further savings in the integral transformation.

With these two approximations, the correlation energy takes the form

$$E_{\text{corr}}^{(2)} = \sum_{(ij) \in P} \frac{2}{1 + \delta_{ij}} \sum_{p,q \in [ij]} K_{pq}^{ij} \tilde{T}_{pq}^{ij}, \qquad (10)$$

where $P$ denotes a list of pairs $(ij)$ with $i \geq j$. Due to its stationary property, the Hylleraas functional computed with the reduced set of amplitudes is still an upper bound to the full MP2 energy. In addition to these reductions of the configuration space, which have been used and tested in previous local correlation methods,[20–22,24] we will investigate in this article two additional approximations of more numerical nature: truncations of the summations in the first and second halves of the integral transformation, eqs. (2) and (6), respectively. These two approximations influence the accuracy of the exchange integrals $K_{pq}^{ij}$, and they may therefore violate the upper bound property of the Hylleraas functional. In fact, as will be discussed in the following sections, the correlation energy is quite sensitive to such approximations and they must therefore be tested carefully. The implementation of these approximations will be described in more detail in the next section.

## Transformation Algorithm

Our local MP2 program[42] is based on the shell-driven integral package of the quantum chemistry program TX96 in which a shell is defined as a set of contracted functions built from the same set of primitive radial Gaussian functions. In Pople's basis sets, $s$ and $p$ functions sharing the same primitive Gaussian exponents are grouped in shells of $L$ (or $sp$) type. In the numerical approximations discussed below whole shells of AO are always considered. This corresponds to the logic of integral computation and guarantees strict rotational invariance. As outlined above, we use four different truncation schemes to reduce the step scaling of the traditional integral transformation algorithms. Those schemes are described in the following. The third and the fourth are new and essential to the present code.

### 1. TRUNCATION OF PROJECTED VIRTUAL SPACE

Substitutions from an orbital pair $(ij)$ are restricted to a pair domain $[ij]$ of projected orbitals $p, q$. The pair domains $[ij]$ are chosen as the direct sum of the *orbital domains* $[i]$ and $[j]$,[21] resulting in square matrices $K^{ij}_{pq}$. An automatic method for the selection of orbital domains was proposed by Boughton and Pulay,[43] and this method has been used in the present work throughout. The only difference from ref. 43 is that we do not restrict the domains to four atoms. While most $\sigma$ bonds and lone pairs are localized on only one or two atoms, in extended $\pi$ systems, like the azo-dye in our test suite, a few orbitals may remain quite delocalized, which consequently leads to larger domains. The selection threshold we used in the Boughton–Pulay algorithm was taken to be the recommended value of 0.02 and was not varied in this work, because it was explored previously by us[21,24] and others.[25,26]

### 2. NEGLECTING DISTANT ORBITAL PAIRS

Correlation between pairs of distant orbitals (weak and distant pairs) can be neglected or calculated approximately. For well-localized orbitals the leading term in the correlation energy between distant orbitals diminishes with the inverse sixth power of the distance between orbital charge centers and is usually below a few $\mu E_h$ if the distance of the orbital centroids exceeds 5 Å. In our older programs, which were directed toward higher level calculations, weak pairs were identified simply by estimating their correlation energy at the lower MP2 level. Because our goal here is to avoid the calculation of MP2 energies, we need a simpler criterion. For well-localized orbitals the distance between their charge centroids is a good indicator of the pair energy and can be used to identify distant pairs. However, for poorly localized orbitals this criterion is not very reliable. Therefore, in all calculations presented in this work, we used the criterion suggested by Hampel and Werner.[24] It is the minimum distance between any of the atoms comprising the orbital domain $[i]$ and those that are part of the domain $[j]$. In the calculations presented here, distant pairs are dropped if the minimum distance exceeds a threshold $T_d$. Neglecting distant pairs contributes much to the efficiency of local MP2 for large systems with respect to memory and computation time requirements.

The question arises whether the effect of neglecting many weak pairs may accumulate to a significant value and may thus lead to inaccurate relative energies. This is possible if two structures are compared that differ greatly in overall shape, such as a compact and an extended structure (e.g., the chain, bowl, and icosahedric structures of $C_{20}$). In most cases, however, the total energy contribution of distant pairs is approximately constant over the energy surface and thus cancels in relative energies. As shown recently,[19] it is possible to model the correlation energies of such pairs empirically; this can account for the bulk of weak pair correlation energy and at least indicate whether weak pair contributions are important. Another possibility is to use multipole expansions, and to approximate the exchange matrices for distant pairs by products of one-electron operator matrix elements. This approach currently is being explored.[18]

### 3. NEGLECTING SMALL CONTRIBUTIONS IN FIRST HALF-TRANSFORMATION

The sparsity of the SCF coefficients can be used to reduce the computational effort in the first half-transformation [eq. (2)]. A shell quartet of integrals is skipped for a given pair $(ij)$ if the product of the maximum integral value $I_{\max}$ and the largest product of MO coefficients involving these shells is smaller than a threshold $T_1$:

$$\frac{|D^{ij}_{\max} I_{\max}|}{C^2_{\mathrm{mean}} P^2_{\mathrm{mean}}} \leq T_1. \qquad (11)$$

$D^{ij}_{\max}$ is the product of the two largest MO coefficients for a given orbital pair $(ij)$ and shell pair, and $C_{\mathrm{mean}}$ and $P_{\mathrm{mean}}$ are mean values of a selected set of MO coefficients of the occupied orbitals and the projected functions, respectively (cf. the Appendix). These were introduced to eliminate basis set dependencies of the criterion caused by $D^{ij}_{\max}$. The Appendix gives a detailed description of this criterion and its rationale. This threshold leads to significant time savings and is essential in our algorithm because it eliminates the formal $O(N^6)$ scaling. It does not affect the memory demand.

## 4. PAIR SPECIFIC TRUNCATION OF VIRTUAL AO SPACE PRIOR TO PROJECTION

Truncations of the summations in the second half-transformation [eq. (6)] limit the AO space in which the exchange matrices $\overline{\mathbf{K}}^{ij}$ have to be computed. This benefits mainly the first half-transformation, which has to be carried out for only a limited number of AOs. Equation (6) is approximated by

$$K^{ij}_{pq} = (pi|jq) = \sum_{\mu\nu\in[ij]_{\mathrm{AO}}} P_{\mu p}\,\overline{K}^{ij}_{\mu\nu}\,P_{\nu q}, \qquad (12)$$

where $[ij]_{\mathrm{AO}}$ denotes an AO domain for pair $(ij)$ (in opposition to the usual pair domain $[ij]$ in projected basis). This truncation reduces the computational effort in the first half-transformation from formally $f\cdot nN^4$ to $f\cdot nN^2L^2$, where $L$ is the average dimension of the AO domains and $f\cdot n$ is the number of orbital pairs $(ij)$. (The remaining factor $N^2$ is reduced by the truncations described in the previous section.) Furthermore, because the indices $\mu, \nu$ of the half-transformed integrals is reduced from $f\cdot nN^2$ to $f\cdot nL^2$. This is essential for the present algorithm because it eliminates the fourth-order dependence of the memory requirement, because $L$ is significantly smaller than the full AO basis for large systems and reaches an asymptotic limit that only depends on the basis set but not on the molecular size. An example of the average size of $L$ in comparison to the full AO basis and the average size of the pair domains in projected AOs has been provided for a sequence of glycine oligomers in Figure 1. Clearly, the accuracy of the correlation energy is sensitive to this truncation, and therefore $L$ must be larger than the domain of pair $ij$ in projected orbitals but becomes constant for extended systems. The reason for this sensitivity is twofold. First, limiting the summations in eq. (6) means approximating the projected orbitals by subsets of AOs, and this may cause a violation of the strong orthogonality condition. Second, the elements of the projection matrix may be large if the projected functions are renormalized, and therefore neglecting small integrals $\overline{K}^{ij}_{\mu\nu}$ may lead to errors in the transformed integrals. As already mentioned, inaccuracies of the transformed integrals cause loss of the variational property of the Hylleraas functional, which is minimized in the iterative local MP2 procedure. The criteria for selecting the truncated AO space must be chosen carefully as shown in our tests.

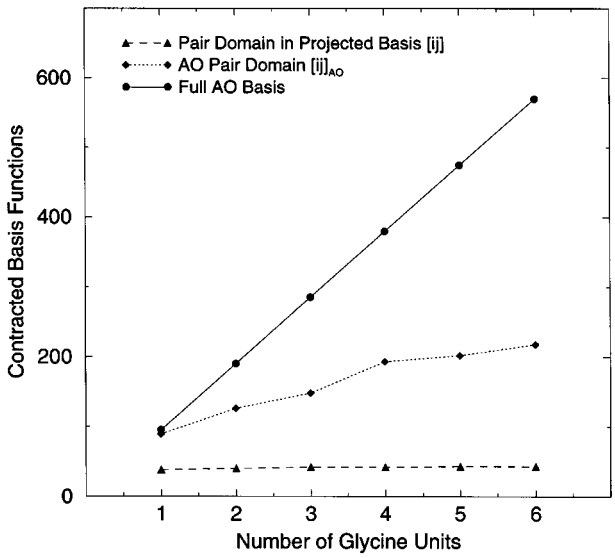


**FIGURE 1.** Comparison of the average size of the pair domains in AOs and projected bases relative to the number of basis functions in an ascending sequence of glycine oligomers. Results are given for a cc-pVDZ basis.

These AO pair domains $[ij]_{\mathrm{AO}}$ are taken to be direct sums of AO orbital domains $[i]_{\mathrm{AO}}$ and $[j]_{\mathrm{AO}}$, which are selected as follows. Our first criterion is based on the orthogonality condition

$$\langle p|i\rangle = (\mathbf{P}^T\overline{\mathbf{S}}\mathbf{C})_{pi} = 0. \qquad (13)$$

The basis functions of a given shell $S$ are included in the AO domain of pair $(ij)$ if

$$\frac{1}{n_S}\sum_{\mu\in S}|(\overline{\mathbf{S}}\mathbf{C})_{\mu i}| \geq T_2, \qquad (14)$$

where $n_S$ is the number of basis functions in the shell and $\overline{\mathbf{S}}$ is the overlap matrix. However, this criterion is not sufficient, because the overlap can be (approximately) zero if the (local) symmetries of the functions $|i\rangle$ and $|p\rangle$ are different. The impact of such cases is minimized by taking the criterion for shells of AOs; if any function in the

shell has the right symmetry, the whole shell is retained. The criterion may still fail, e.g., for the overlap of a $\pi$-type MO with an $s$ function. Ideally, we would use the *absolute* overlap $\langle|\mu||i|\rangle$ between a localized orbital and a basis function, but this is difficult to calculate. It can be simulated, however, by the additional condition

$$\left[\frac{v}{n_S}\sum_{\mu\in S}\sqrt{X_{\mu i}^2+Y_{\mu i}^2+Z_{\mu i}^2}\geq T_2\right], \qquad (15)$$

where, for example,

$$X_{\mu i}=\langle\mu|\hat{x}-x^{(i)}|i\rangle=(\overline{\mathbf{X}}\mathbf{C})_{\mu i}-x^{(i)}(\overline{\mathbf{S}}\mathbf{C})_{\mu i}, \qquad (16)$$

and $Y_{\mu i}$ and $Z_{\mu i}$ are defined analogously. Here $\hat{x}$, $\hat{y}$, and $\hat{z}$ are the coordinate operators; $x^{(i)}$, $y^{(i)}$, $z^{(i)}$ are the coordinates of the charge center of the localized MO $i$; and $v$ is a scaling factor that was empirically taken to be 0.12 $\mathrm{au}^{-1}$ without further refinement. These dipole integrals give nonzero values for a $\sigma$ AO and a $\sigma$ MO if they spatially overlap. We have found it advantageous to use a lower thresholds $T_2$ (by about a half-order of magnitude) for diagonal and close orbital pairs (i.e., pairs that share one atom), while all other pairs can be treated slightly less accurately. Poorly localized orbitals in large $\pi$ systems also require a sharper threshold. Useful ranges for threshold $T_2$ will be discussed later.

As is clear from the description of these criteria, they leave space for experimentation and further improvements. Even though these thresholds needed adjustment for some critical cases as discussed below, they work in most cases. The whole transformation can be summarized as follows:

1. First we create a list of correlated orbital pairs $(ij)$. Distant pairs are neglected according to threshold $T_d$ as described above. This list is further reduced by other criteria in the subsequent steps.
2. The integrals in a given shell quartet $(MR|SN)$ are sorted into matrices $(\mathbf{I}')_{\mu'\nu',\rho'\sigma'}=(\mu\rho|\sigma\nu)$, where Greek indices denote individual AOs and capital letters refer to shells ($\mu\in M$, $\nu\in N$, etc.) Here and in the following the primes indicates quantities whose dimensions are determined by the shell sizes. The absolute AO indices $\mu$ are related to the shell indices $\mu'$ by shell offsets (i.e., the absolute address of the first basis function in shell $\mu'$).
3. Permutational symmetry is used fully (i.e., each integral is used eight times), transforming the indices belonging to the shells RS, RN, MS, and MN. Each of these cases gives two contributions (supermatrix symmetry). In the following we outline the algorithm only for one of the cases, the one in which the shells RS are transformed to occupied orbitals. The other cases are entirely similar.
4. We generate a list of all nondistant pairs $(ij)$ according to the following criteria: (1) threshold $T_1$ [cf. eq. (11)] is exceeded for the shells $R$, $S$; and (2) threshold $T_2$ [cf. eqs. (14) and (15)] is exceeded for the shells $M$, $N$. For this reduced list of pairs we generate the matrices

$$(\mathbf{D}')_{\rho'\sigma',ij}=C_{\rho' i}C_{\sigma' j}, \qquad (17)$$

$$(\mathbf{D}')_{\rho'\sigma',ji}=C_{\rho' j}C_{\sigma' i}. \qquad (18)$$

   Here $ij$ and $ji$ are to be considered as indices of a reduced pair list.
5. The actual first half-transformation [cf. eq. (2)] uses efficient matrix multiplications

$$\overline{\mathbf{K}}'=\mathbf{I}'\cdot\mathbf{D}'. \qquad (19)$$

   Shells with high angular momentum functions give rise to larger blocks and therefore increase the efficiency of the matrix multiplication.
6. In the last step the elements of the small matrices $\mathbf{K}'$, whose dimension is determined by the sizes of shells $M$ and $N$, are scattered to their final locations in the internal exchange matrices $\overline{\mathbf{K}}^{ij}$.

$$(\overline{\mathbf{K}}^{ij})_{\mu\nu}=(\overline{\mathbf{K}}^{ij})_{\mu\nu}+(\overline{\mathbf{K}}')_{\mu'\nu',ij}, \qquad (20)$$

$$(\overline{\mathbf{K}}^{ij})_{\nu\mu}=(\overline{\mathbf{K}}^{ij})_{\nu\mu}+(\overline{\mathbf{K}}')_{\mu'\nu',ji}. \qquad (21)$$

After all integral shell quartets have been processed, the exchange matrices are transformed into the projected basis according to eq. (6). This requires just two matrix multiplications per pair, and the time for this step is quite negligible. Subsequently, the MP2 equations are solved iteratively and the energy is computed.

Essentially all of the logic is required to generate the effective pair list and to address the local internal exchange matrices. The transformation of the shell blocks is free of time consuming logic. If there is sufficient memory to store all local exchange matrices $\overline{\mathbf{K}}^{ij}$, then this algorithm requires only a single pass through the integrals. However,

if this condition is not met, then it is possible to calculate the internal exchange matrices for a subset of orbital pairs by projecting them and storing the resulting small local exchange matrices on disk. This algorithm is essentially open ended but requires the repeated evaluation of the integrals for each batch of pairs. Parallelizations of the algorithm appears feasible and will be explored in the future.

## Effects of Thresholds

This section discusses the effects of the four truncation schemes on the results. The simplest is threshold $T_d$ (i.e., the neglect of pair correlation energies for distant pairs). The loss of correlation energy caused by this grows with the number of pairs neglected. The appropriate value of $T_d$ can be estimated by using averaged pair energies for various distance ranges. Because of these simple characteristics, and because such tests have been performed in previous work, we did not test this truncation explicitly. Unless otherwise noted, we have used a threshold of 6.0 Å for all molecules tested, which appears to be a safe choice. The Pipek–Mezey localization scheme[44] has been used in all calculations, even though it leads to larger pair domains for extended $\pi$-systems than the Boys localization.[45]

Truncation schemes $T_1$ and $T_2$ have been tested with several basis sets for lactic acid (18 correlated orbitals) and glycine (15 correlated orbitals). The results are summarized in Table I. Two families of basis sets have been used; Dunning's recent cc-pVDZ and cc-pVTZ bases,[46] and 6–31G and 6–311G sets augmented with polarization functions.[47] The first family is generally contracted while the second has segmented contraction.

The error in the correlation energy caused by threshold $T_1$ (truncations in the first half transformation), shown in column 8 of Table I, is small (around 0.01%) for most basis sets. However, for the cc-pVDZ we find large errors of up to 0.09%. This larger error is caused by a steep increase of the average magnitude of the MO coefficients for small bases, leading to a failure of the correction term for basis set independency [i.e., the denominator in eq. (11)]. Refining this empirical correction should resolve this problem. However, the error is basis set specific and vanishes for larger bases (cc-pVTZ). Particularly critical are basis sets containing diffuse functions. Tests on 6–31+G($d$), 6–311++G($d$, $p$), and aug-cc-pVDZ bases have shown that criterion $T_1$ is not yet able to handle these functions properly. Considering a 6–311++G($d$, $p$) calculation on lactic acid, $T_1$ introduces an error of 0.335% of the total correlation energy; $T_2$ leads to an inaccuracy of only 0.002%, which is in the usual range as shown in Table I (thresholds for this calculation are the same as given in Table I). Thus, using this product criterion is not yet recommended for basis sets containing diffuse functions.

Threshold $T_2$ (truncation of the virtual AO space prior to projection) appears to be less basis set dependent than $T_1$. Generally contracted bases perform slightly worse than segmented functions, but the former behave more systematically. For most basis sets tested the numerical error introduced by threshold $T_2$ is less than 0.01% of the total correlation energy. The combination of thresholds $T_1$ and $T_2$ leads to about a 0.02% error in the correlation energy (with the exceptions discussed above). For all molecules and basis sets tested, using thresholds $T_1$ and $T_2$ yielded correlation energies that are too low. This is characteristic of truncations that violate the variation theorem and leads to a fortuitous error cancellation when invoking threshold $T_d$ for distant pairs, as seen in the benchmark calculations below.

Stability with respect to molecular size has been tested using the 6–31G($d$) basis and a set of eight test molecules given in Table II. The molecules chosen show a variety of different functional groups and $\pi$ systems of varying size. The average numerical error for all molecules tested is less than 0.01% for the parameter values used. We could not find systematic tendencies with respect to the molecules size for either $T_1$ or $T_2$. Individual characteristics of the molecules seem to have a much stronger impact on the quality of the results than just the molecular size. The somewhat larger error for coffein introduced by $T_2$ can probably be attributed to poor localization of the $\pi$ system, which indeed leads to pair domains [$ij$] up to 10 (!) atoms. In most cases pair domains include less than 6 atoms. However, the azo-dye[51] given in Table II also has a very delocalized $\pi$ system with pair domains up to 11 atoms and shows much smaller errors. Systems like these are stringent tests for the AO domain selection and certainly show the limits of our proposed criterion. Even though one can eliminate the error in coffein easily by using a lower threshold, it would be desirable to change the criterion in such a way that a con-

**TABLE I.**
**LMP2 Correlation Energies and Numerical Errors (in $10^{-2}$%) as Introduced by Thresholds Dependent on Basis Set Used.**

| | | | Test for $T_2$ | | Test for $T_1$ | | | |
|---|---|---|---|---|---|---|---|---|
| | | $E_{LMP2}^{Ref.}$ | $E_{LMP2}$ | | $E_{LMP2}$ | | | |
| Basis Set | CF[a] No. | $T_1 = 0.000$ $T_2 = 0.000$ | $T_1 = 0.000$ $T_2 = 0.004$ | Error[b] | $T_1 = 9.10^3$ $T_2 = 0.004$ | Error[b] | Error[c] | Ref. |
| | | | Lactic Acid | | | | | |
| 6-31G(*d*) | 102 | −0.89899090 | −0.89899182 | 0.010 | −0.89903361 | 0.475 | 0.476 | 47 |
| 6-31G(*d*, *p*) | 120 | −0.93234490 | −0.93244470 | 1.070 | −0.93247581 | 1.404 | 0.334 | 47 |
| 6-311G(*d*, *p*) | 144 | −1.00426821 | −1.00426179 | 0.064 | −1.00428220 | 0.139 | 0.203 | 48 |
| 6-311G(2*d*, *p*) | 174 | −1.06923087 | −1.06928618 | 0.517 | −1.06933294 | 0.955 | 0.437 | 48 |
| 6-311G(2*d*, 2*p*) | 192 | −1.07547426 | −1.07553284 | 0.545 | −1.07557146 | 0.904 | 0.359 | 48 |
| 6-311G(2*df*, 2*pd*) | 264 | −1.17906340 | −1.17922405 | 1.363 | −1.17926544 | 1.714 | 0.351 | 48 |
| cc-pVDZ[3*s*2*p*1*d* / 2*s*] | 96 | −0.90681822 | −0.90685535 | 0.409 | −0.90761378 | 8.773 | 8.363 | 49 |
| cc-pVDZ[3*s*2*p*1*d* / 2*s*1*p*] | 114 | −0.94204484 | −0.94207054 | 0.273 | −0.94252658 | 5.114 | 4.841 | 49 |
| cc-pVTZ[4*s*3*p*2*d* / 3*s*2*p*] | 192 | −1.08878731 | −1.08884136 | 0.496 | −1.08888215 | 0.871 | 0.375 | 49 |
| cc-pVTZ[4*s*3*p*2*d*1*f* / 3*s*2*p*] | 234 | −1.18243173 | −1.18249497 | 0.543 | −1.18252672 | 0.803 | 0.260 | 49 |
| cc-pVTZ[4*s*3*p*2*d*1*f* / 3*s*2*p*1*d*] | 264 | −1.19065272 | −1.19071300 | 0.506 | −1.19075638 | 0.871 | 0.364 | 49 |
| | | | Glycine | | | | | |
| 6-31G(*d*) | 85 | −0.77338040 | −0.77339406 | 0.177 | −0.77340267 | 0.288 | 0.111 | 47 |
| 6-31G(*d*, *p*) | 100 | −0.79907901 | −0.79909218 | 0.165 | −0.79910863 | 0.371 | 0.206 | 47 |
| 6-311G(*d*, *p*) | 120 | −0.85487988 | −0.85487783 | 0.024 | −0.85489175 | 0.139 | 0.163 | 48 |
| 6-311G(2*d*, *p*) | 145 | −0.90847054 | −0.90848927 | 0.206 | −0.90853979 | 0.762 | 0.556 | 48 |
| 6-311G(2*d*, 2*p*) | 160 | −0.91184823 | −0.91186435 | 0.177 | −0.91195605 | 1.182 | 1.006 | 48 |
| 6-311G(2*df*, 2*pd*) | 220 | −0.99718143 | −0.99720337 | 0.220 | −0.99736787 | 1.870 | 1.650 | 48 |
| cc-pVDZ[3*s*2*p*1*d* / 2*s*] | 80 | −0.77874478 | −0.77881010 | 0.839 | −0.77942891 | 8.785 | 7.946 | 49 |
| cc-pVDZ[3*s*2*p*1*d* / 2*s*1*p*] | 95 | −0.80348198 | −0.80353506 | 0.661 | −0.80394824 | 5.803 | 5.142 | 49 |
| cc-pVTZ[4*s*3*p*2*d* / 3*s*2*p*] | 160 | −0.92426394 | −0.92430537 | 0.448 | −0.92441831 | 1.670 | 1.222 | 49 |
| cc-pVTZ[4*s*3*p*2*d*1*f* / 3*s*2*p*] | 195 | −1.00200186 | −1.00205215 | 0.502 | −1.00217422 | 1.720 | 1.218 | 49 |
| cc-pVTZ[4*s*3*p*2*d*1*f* / 3*s*2*p*1*d*] | 220 | −1.00836828 | −1.00841593 | 0.473 | −1.00855113 | 1.813 | 1.341 | 49 |

Distant pairs have not been neglected, and the criterion of 0.02 for the virtual space selection as suggested by Boughton and Pulay has been used throughout.[43]

[a] Number of contracted functions.

[b] Relative to column 3.

[c] Relative to column 4.

stant threshold would work in all cases. On the basis of our tests, we recommend a threshold of about $T_1 = 5 \cdot 10^3$ for the product criterion and $T_2 = 0.004$ for the AO space truncation. For generally contracted basis sets, somewhat sharper thresholds should be chosen.

## Benchmark Calculations

For testing the performance of our code, we chose a sequence of glycine oligomers up to the hexamer ($[C_3H_5NO_2]_x$; $x = 1$–$6$). The geometries were taken from an extended crystallographic unit cell[52] (Fig. 2). This sequence shows the scaling of the method directly, because the number of correlated orbitals and the number of contracted functions increase linearly with the number of monomer units. It is also useful to explore the effect of molecular shape on the scaling. Up to the trimer the shape of the cluster is essentially 2-dimensional (2-D), while higher oligomers are 3-D. Lower dimension (linear and planar) structures are expected to scale better than 3-D ones.

Table III shows the memory requirements and CPU times for two different basis sets: Dunning's cc-pVDZ and cc-pVTZ bases,[46] the latter without *d* functions on the hydrogens. Two sets of thresholds

**TABLE II.**
**Correlation Energies and Numerical Errors (in $10^{-2}$%) as Introduced by Thresholds Dependent on the Molecular Size.**

| | | | | Test for $T_2$ | | Test for $T_3$ | | |
|---|---|---|---|---|---|---|---|---|
| | | | $E^{\text{Ref.}}_{\text{LMP2}}$ | $E_{\text{LMP2}}$ | | $E_{\text{LMP2}}$ | | |
| Molecule | No. CF[a] | No. CE[b] | $T_1 = 0.000$ $T_2 = 0.000$ | $T_1 = 0.000$ $T_2 = 0.004$ | Error[c] | $T_1 = 4 \cdot 10^3$ $T_2 = 0.004$ | Error[c] | Error[d] |
| Acetone | 72 | 24 | −0.55480941 | −0.55481146 | 0.037 | −0.55481207 | 0.048 | 0.011 |
| Glycine | 85 | 30 | −0.77338040 | −0.77339406 | 0.177 | −0.77339973 | 0.250 | 0.073 |
| Lactic acid | 102 | 36 | −0.89899090 | −0.89899182 | 0.010 | −0.89902119 | 0.337 | 0.327 |
| Uracil | 128 | 42 | −1.15300124 | −1.15301682 | 0.135 | −1.15307166 | 0.611 | 0.476 |
| Adenine | 160 | 50 | −1.41041784 | −1.41049751 | 0.565 | −1.41054240 | 0.883 | 0.318 |
| Glycine dimer | 170 | 60 | −1.54524311 | −1.54527699 | 0.219 | −1.54528279 | 0.257 | 0.038 |
| Menthol | 205 | 66 | −1.45089391 | −1.45089604 | 0.015 | −1.45094701 | 0.366 | 0.351 |
| Coffein | 230 | 74 | −1.97181493 | −1.97208814 | 1.386 | −1.97222908 | 1.647 | 0.715 |
| HPAMP[e] | 262 | 80 | −2.12201892 | −2.12203600 | 0.080 | −2.12211215 | 0.439 | 0.359 |

All calculations refer to a 6–31G($d$) basis. Distant pairs have not been neglected, and the criterion of 0.02 for the virtual space selection as suggested by Boughton and Pulay has been used throughout.[43]
[a] Number of contracted functions.
[b] Number of correlated electrons.
[c] Relative to column 4.
[d] Relative to column 5.
[e] (4-[(4′-hydroxyphenyl)azo]-*N*-methylpyridine.

have been chosen for the cc-pVDZ calculations to show the variation of computing times and memory requirements with the accuracy of the calculations. The average dimension of the AO domains $[ij]_{\text{AO}}$ is also given. Even though the AO domains are quite large, they grow much slower than the number of basis functions (cf. Fig. 1). For the double-zeta basis chosen, they probably saturate below 300 functions for most molecules. For the more diffuse cc-pVTZ bases, the limit is higher. It is clear that our largest systems are still not large enough to reach the limiting efficiency of the local MP2 method. For larger calculations, which exceed our present computing resources but which will certainly be feasible soon, the advantage of our local method over traditional MP2 should be more pronounced.

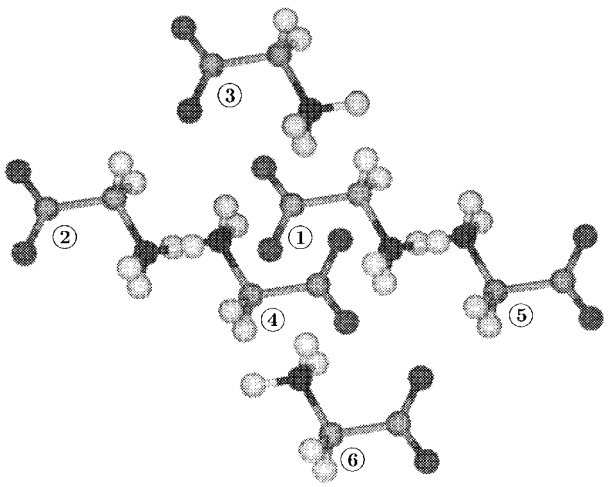


**FIGURE 2.** Structure of the glycine hexamer as taken from its crystallographic unit cell. The numbers are the unit numbers (i.e., a glycine trimer would be built from units 1, 2, and 3).

The percentage of neglected distant pairs is shown in column 6 of Table III. For the largest system, the glycine hexamer, only 42% of the pairs are neglected. We expect that we will be able to improve the code significantly by using approximate methods for the calculation of distant but not negligible pairs.[18] The percentage of neglected pairs clearly shows the influence of molecular shape: in going from the essentially planar trimer to the 3-dimensional tetramer, the percentage decreases.

The memory requirements scales with about the 2.5th power of the molecular size for calculations using one integral pass. This is a significant

**TABLE III.**
**Benchmark Calculations for Sequence of Glycine Oligomers with Structures Taken from Crystallographic Data.**

| Molecule | Basis[a] | CF[b] | AV[c] No. | CE[d] | NP[e] | Memory Trans. | Memory LMP2 | Memory Disk | CPU time Trans. | CPU time ERIs | CPU time Total | Correlation energies $E_{LMP2}$ | Correlation energies $E_{Ref}$ |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Monomer | cc-pVDZ | 95 | 90 | 30 | 0.0 | 1.3 | 0.7 | 0.0 | 6 | 1 | 8 | −0.80369515 | −0.80348228 |
| | | | 89 | | | 1.3 | | | 6 | 1 | 8 | −0.80379915 | |
| Dimer | cc-pVDZ | 190 | 130 | 60 | 17.2 | 8.0 | 2.5 | 0.0 | 50 | 4 | 56 | −1.60590621 | −1.60577350 |
| | | | 126 | | | 7.5 | | | 43 | 4 | 49 | −1.60599622 | |
| Trimer | cc-pVDZ | 285 | 154 | 90 | 27.7 | 21.1 | 5.3 | 0.0 | 154 | 10 | 171 | −2.41633266 | −2.41637237 |
| | | | 148 | | | 19.6 | | | 131 | 10 | 141 | −2.41643988 | |
| Tetramer | cc-pVDZ | 380 | 203 | 120 | 24.2 | 64.5 | 9.8 | 0.0 | 753 | 25 | 801 | −3.22048798 | −3.22062184 |
| | | | 193 | | | 58.6 | | | 612 | 25 | 661 | −3.22056705 | |
| Pentamer | cc-pVDZ | 475 | 215 | 150 | 36.5 | 92.0 | 13.2 | 5.7 | 1408 | 42 | 1495 | −4.02948436 | — |
| | | | 202 | | | 81.3 | | | 996 | 42 | 1083 | −4.02957512 | |
| Hexamer[f] | cc-pVDZ | 570 | 236 | 180 | 42.1 | 68.5 | 8.6 | 23.2 | 2306 | 131 | 2538 | −4.84085174 | — |
| | | | 218 | | | 68.5 | | | 1633 | 129 | 1863 | −4.84089500 | |
| Monomer | cc-pVTZ | 195 | 184 | 30 | 0.0 | 5.5 | 2.9 | 0.0 | 65 | 9 | 77 | −1.00219536 | −1.00200204 |
| Dimer | cc-pVTZ | 390 | 267 | 60 | 17.4 | 32.8 | 9.8 | 0.0 | 616 | 34 | 668 | −2.00448738 | −2.00441442 |
| Trimer | cc-pVTZ | 585 | 322 | 90 | 27.8 | 84.5 | 10.8 | 20.8 | 2397 | 87 | 2549 | −3.01566041 | — |
| Tetramer[g] | cc-pVTZ | 780 | 424 | 120 | 24.2 | 78.5 | 16.5 | 39.9 | 11790 | 836 | 12828 | −4.02324470 | — |
| Pentamer[h] | cc-pVTZ | 975 | 454 | 150 | 36.6 | 90.5 | 22.7 | 52.1 | 20639 | 2064 | 23118 | −5.03310598 | — |

CPU times are given in minutes and refer to an SGI Power Challenge R10000 / 194 MHz with 1.5 GB memory and 2 MB secondary cache. Memory and disk space requirements are given in MW, correlation energies in $E_h$. Thresholds used for the first set of cc-pVDZ calculations: $T_d = 6.0$ Å, $T_2 = 0.005$, $T_1 = 10^3$; for the second set of cc-pVDZ calculations: $T_d = 6.0$ Å, $T_2 = 0.006$, $T_1 = 4 \cdot 10^3$. Thresholds used for cc-pVTZ benchmarks: $T_d = 6.0$ Å, $T_2 = 0.005$, $T_1 = 9 \cdot 10^3$. The criterion of 0.02 for the virtual space selection as suggested by Boughton and Pulay has been used throughout.[43]

[a] Dunning's cc-pVTZ basis has been truncated for hydrogens [4*s*3*p*2*d*1*f* / 3*s*2*p*].

[b] Number of contracted functions.

[c] Average dimension of an AO pair domain $[ij]_{AO}$.

[d] Number of correlated electrons.

[e] Percentage of neglected orbital pairs.

[f] Two integral passes.

[g] Four integral passes.

[h] Six integral passes.

reduction relative to the $n^2N^2$ scaling for full conventional in-core transformations. Disk space requirements are easily satisfied by today's large (multi-gigabyte) disks. Loosening the thresholds as given in the lower set of numbers of Table III leads to about a 10% memory savings. The bottleneck with respect to memory requirements is still the transformation step, because the AO pair domains $[ij]_{AO}$ have to be significantly larger than the $[ij]$ domains in projected space. Savings in CPU time due to looser thresholds amount to about 30%. However, the overall CPU scaling of about $O(N^{2.6})$ to $O(N^{2.9})$ is not affected significantly by choosing different parameters. This is evident in Figure 3 that shows the CPU times of the integral transformation as a function of the number of glycine monomers for two choices of thresholds. The relatively steep increase in CPU time between the trimer and the tetramer corresponds to the transition from 2-D to 3-D clusters. Estimating the CPU time for a full SCF calculation as 10 full integral evaluations, a local MP2 calculation takes roughly twice as long as the preceding SCF step. The results for cc-pVTZ calculations resemble those for cc-pVDZ benchmarks. The fraction of CPU time spent in the integral calculation becomes more significant for the cc-pVTZ benchmarks. This is due to the memory limitations that require several integral passes for larger clusters. The efficiency of the thresholds is shown by the reduction of the formal $O(N^6)$ scaling to the observed $\approx O(N^3)$ or lower [$O(N^{2.5})$ in the best case]. The benign scaling of our method makes it feasible to treat large molecules, which could so far be treated at the SCF or density functional theory level only, with MP2 on single-processor workstations. However, the rather large prefactor caused by the logic, as well as the scatter and gather operations, makes the method only competitive for large systems. The crossover point relative to conventional MP2

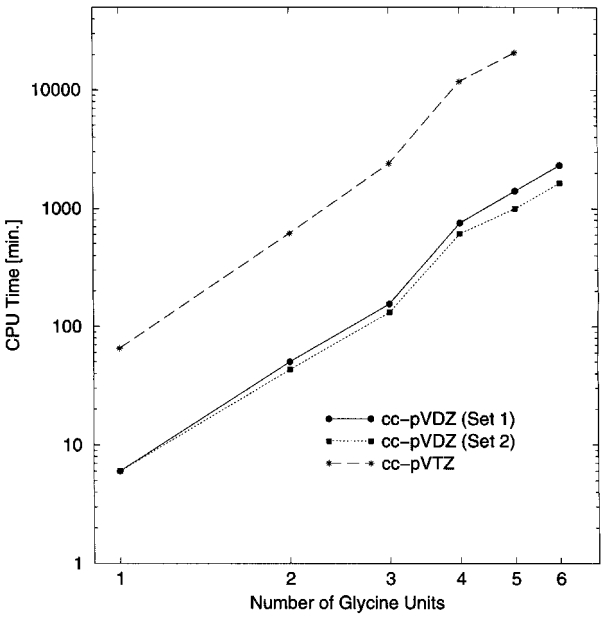


**FIGURE 3.** Performance of the transformation step for a sequence of glycine units. The set labeling and the values presented refer to Table III.

should be around 400 basis functions for cc-pVDZ basis sets. For those oligomers for which local reference calculations could be performed (see Table III), the numerical error of the total correlation energy varies from 0.073 to 0.317 $mE_h$. For small clusters the absolute value of the correlation energy is too high; it is too low for larger clusters. This effect is probably caused by neglecting distant orbital pairs that accumulate for larger systems and could be eliminated by an approximate treatment of long-range correlation. Because the number of neglected pairs is related to the shape of the molecule (see above), the deviations caused by this threshold are as well.

## Conclusions

A single step two index transformation of electron repulsion integrals to localized MOs offers, despite a formally unfavorable $O(N^6)$ scaling, two advantages. First, it allows an effective use of prescreening techniques, which dramatically reduce the number of arithmetic operations and thus the scaling to about $O(N^3)$. Second, because the number of half-transformed integrals is less than 15% of the value in the traditional method, the bottleneck of storing huge arrays of transformed integrals is largely eliminated. The algorithm presently allows a memory efficient treatment of molecules up to 1000 basis functions on customary workstations, opening the use of more correlation methods for larger molecules than previously possible. The largest savings are expected for MP2, but all correlation methods could profit from this new transformation. A disadvantage is that it requires the orbital-invariant formulation of MP2, which is somewhat more costly than the canonical formation. The actual scaling of our local MP2 method is slightly lower than $\sim N^3$, making this approach competitive with other nontraditional methods, for instance the pseudospectral LMP2 of Murphy et al.[25] Compared to traditional MP2, the method offers significant advantages only for large systems with around 400 basis functions or more. However, because LMP2 calculations needing more than 600 basis functions are still quite expensive when using the algorithm presented, a small window remains between 400 and 600 basis functions for routine applications at present. But, with a moderate future perspective of an increase in floating point operations by a factor of 10 within the next 5 years, our largest calculation can be performed as a standard application easily by then. Moreover, making use of modern multiprocessor workstations or ven large MPP supercomputers for this algorithm, which is currently being explored by one of us, will lead to an additional speedup and therefore we anticipate routine applications of up to more than 1500 basis functions soon.

The use of truncations introduces a small error of about 0.01% of the correlation energy relative to a local MP2 calculation with full numerical accuracy. For very large molecules, this amount may come close to the accuracy desired for reaction paths or rotational barriers. For example, if the total correlation energy amounts to 10 $E_h$, the error may exceed 0.5 kcal/mol. However, it is expected that much of this error largely cancels out in calculating relative energies.[53] At any rate, it can be reduced easily by using tighter thresholds if higher accuracy is needed. Friesner's analytically corrected pseudospectral LMP2 also shows numerical deviations in the same range [reported to be typically within 0.2 kcal/mol, e.g., 0.09 kcal/mol (6–31G**) for the glycine moner].[25] Weigend and Häser recently presented a systematic study on RI-MP2 energies for a large set of molecules.[37] For large systems, as considered in this study, they obtain deviations in the range of about 0.05%

relative to MP2 correlation energies [benzene, 0.01% (0.07 kcal/mol); porphyrin, 0.07% (1.45 kcal/mol)], but they are also able to show that most of these deviations cancel out when considering relative reaction energies. We expect the same for the local approach.

All these absolute errors are insignificant compared to systematic errors in conventional MP2 calculations, which are caused by the MP2 approximation itself, basis set superposition errors, and the fact that for medium size basis sets typically only ~80% of the MP2 correlation energy will be recovered. Further refinements of the local method can be expected from using *mathematically* localized orbitals, which minimize the number of significant SCF coefficients, and the approximate treatment of distant pairs. Moreover, the rather large prefactor, which is responsible for placing the crossover point to about 400 basis functions and 1000 pairs (cc-pVDZ), should decrease by a better addressing scheme that takes into account the size of the primary cache in dealing with the compressed internal exchange matrices.

## Appendix: Derivation for Criterion to Exploit MO Coefficient Sparsity

The threshold introduced that neglects insignificant coefficient products ($T_1$) acts on a shell level. We have generated a matrix (shells × valence orbitals) that contains the largest absolute LMO coefficient per shell and orbital, $C_{Ri}^{\max}$. Likewise, we search for the largest absolute electron repulsion integral within a given shell quartet. Therefore, early after obtaining shell indices for a block of two electron integrals, we can decide about the significance of contributions to internal exchange matrices according to

$$|D_{\max}^{ij} I_{\max}| = |C_{Ri}^{\max} C_{Sj}^{\max} (MR|SN)^{\max}| \le T_1^{\dagger}$$
$$M, N, R, S = \text{shell indices.} \quad \text{(A.1)}$$

A useful range for threshold $T_1^{\dagger}$ (which is different in magnitude than the finally used threshold $T_1$) is $10^{-7}$. Unfortunately, this simple criterion appears to be basis set dependent, because the average magnitude of the MO coefficients strongly depends on the basis set chosen. In order to get a basis set independent criterion, the expression in eq. (A.1) has been divided by the square of the average magnitude of the largest shell coefficients of the localized MOs and of the projected AOs:

$$\frac{|D_{\max}^{ij} I_{\max}|}{C_{\text{mean}}^2 P_{\text{mean}}^2} \le T_1 \quad \text{(A.2)}$$

where

$$C_{\text{mean}} = \frac{1}{n_{Tk}} \sum_k \sum_{T \in |k|_{\text{AO}}} |C_{Tk}^{\max}| \quad \text{(A.3)}$$

is the average of all $C_{Tk}^{\max}$ with the restriction that the AO shell $T$ contributes to MO domain $[k]_{\text{AO}}$.

$$P_{\text{mean}} = \frac{1}{n_{Tq}} \sum_k \sum_{T,q \in [k]_{\text{AO}}} |P_{Tq}^{\max}| \quad \text{(A.4)}$$

is in analogy the average of all $P_{Tq}^{\max}$ (being defined as $C_{Tk}^{\max}$) where $T$ comprises all AO shells of orbital domain $[k]_{\text{AO}}$ and $q$ runs over all projected functions within this domain. Strictly, one would have to run $q$ over all elements in $[k]$, but this leads to a too small selection of significant elements of the projection matrix. This correction of eq. (A.1) leads to a coupling of the criterion for the virtual space selection and the criterion for small contributions. This coupling is necessary to restrict the averaging to significant coefficients, because the bulk of coefficients is negligible for localized MOs.

## Acknowledgments

This work was supported by the Deutsche Forschungsgemeinschaft (DFG) in the Schwerpunktsprogramm Molekulare Cluster. It was also supported in part by a grant of the Air Force Office of Scientific Research to P. P. and the National Science Foundation under Grant CHE-9319929. P. P. thanks the Alexander von Humboldt Foundation for a Senior Scientist Award, the National Science Foundation for support under Grant CHE-9707202, and the Air Force Office for support under Grant F49620-94-1. H. J. W. acknowledges the EU support in the TMR network FMRX-CT96-088 and the Fonds der Chemischen Industrie.

## References

1. K. C. Tang and C. Edmiston, *J. Chem. Phys.*, **52**, 997 (1970).
2. S. T. Elbert, In *Numerical Algorithms in Chemistry: Algebraic Methods*, Reports of NRCC Workshop 1978, LBL-8158, 1978, p. 128.
3. M. Yoshimine, Report RJ-555 IBM Research Laboratory, San Jose, CA, 1969.

4. H. J. Werner and W. Meyer, *J. Chem. Phys.*, **73**, 2342 (1980).
5. V. R. Saunders and J. H. van Lenthe, *Mol. Phys.*, **48**, 923 (1983).
6. P. R. Taylor, *Int. J. Quantum Chem.*, **31**, 521 (1987).
7. S. Saebø and J. Almlöf, *Chem. Phys. Lett.*, **154**, 83 (1989).
8. M. Head–Gordon, J. A. Pople, and M. J. Frisch, *Chem. Phys. Lett.*, **153**, 503 (1988).
9. M. J. Frisch, M. Head–Gordon, and J. A. Pople, *Chem. Phys. Lett.*, **166**, 281 (1990).
10. S. Wilson, In *Methods in Computational Chemistry*, S. Wilson, Ed., Plenum Press, New York, 1987, p. 251.
11. I. Shavitt, In *Methods of Molecular Electronic Structure Theory*, H. F. Schaefer, ed., Plenum Press, New York, 1977, p. 189.
12. A. M. Márquez and M. Dupuis, *J. Comput. Chem.*, **16**, 395 (1995).
13. I. M. B Nielsen and E. T. Seidl, *J. Comput. Chem.*, **16**, 1301 (1995).
14. D. E. Bernholdt and R. J. Harrison, *J. Comput. Chem.*, **102**, 9582 (1995).
15. A. T. Wong, R. J. Harrison, and A. P. Rendell, *Theor. Chim. Acta*, **93**, 317 (1996).
16. M. Schütz and R. Lindh, *Theor. Chim. Acta*. **95**, 13 (1997).
17. M. Schütz, R. Lindh, and H.-J. Werner, unpublished data.
18. G. Hetzer, P. Pulay, and H.-J Werner, *Chem. Phys. Lett.*, in press.
19. G. Rauhut, J. W. Boughton, and P. Pulay, *J. Chem. Phys.*, **103**, 5662 (1995).
20. P. Pulay and S. Seabø, *Theor. Chim. Acta*, **69**, 357 (1986).
21. (a) S. Seabø and P. Pulay, *J. Chem. Phys.*, **86**, 914 (1987); (b) S. Saebø and P. Pulay, *J. Chem. Phys.*, **88**, 1884 (1988).
22. S. Seabø and P. Pulay, *Annu. Rev. Phys. Chem.*, **44**, 213 (1993).
23. J. Almlöf, K. Faegri, and K. Korsell, *J. Comput. Chem.*, **3**, 385 (1982).
24. C. Hampel and H.-J. Werner, *J. Chem. Phys.*, **104**, 6286 (1996).
25. R. B. Murphy, M. D. Beachy, R. A. Friesner, and M. N. Rignalda, *J. Chem. Phys.*, **103**, 1481 (1995).
26. G. Reynolds, T. J. Martinez, and E. A. Carter, *J. Chem. Phys.*, **105**, 6455 (1996).
27. A. El Azhary, G. Rauhut, P. Pulay, and H.-J. Wener, *J. Chem. Phys.*, **108**, 5185 (1998).
28. R. Knab, W. Förner, J. Cizek, and J. Ladik, *J. Mol. Struct. (Theochem.)*, **366**, 11 (1996).
29. E. Kapuy and C. Kozmutza, *J. Chem. Phys.*, **94**, 5565 (1991).
30. K. Szalewicz, B. Jeziorski, H. J. Monkhorst, and J. G. Zabolitzky, *Chem. Phys. Lett.*, **91**, 169 (1982).
31. M. Häser and J. Almlöf, *J. Chem. Phys.*, **96**, 489 (1992).
32. G. Rauhut and P. Pulay, *Chem. Phys. Lett.*, **248**, 223 (1996).
33. A. K. Wilson and J. Almlöf, *Theor. Chim. Acta*, **95**, 49 (1997).
34. A. Komornicki and G. Fitzgerald, *J. Chem. Phys.*, **98**, 1398 (1993).
35. M. Feyereisen, G. Fitzgerald, and A. Komornicki, *Chem. Phys. Lett.*, **208**, 359 (1993).
36. D. E. Bernholdt and R. J. Harrison, *Chem. Phys. Lett.*, **250**, 477 (1996).
37. F. Weigend and M. Häser, *Theor. Chem. Acc.*, **97**, 331 (1997).
38. W. Meyer, *J. Chem. Phys.*, **64**, 2901 (1976).
39. P. Pulay, S. Saebø, and W. Meyer, *J. Chem. Phys.*, **81**, 1901 (1984).
40. P. Pulay, *Chem. Phys. Lett.*, **73**, 393 (1980).
41. S. Saebø, W. Tong, and P. Pulay, *J. Chem. Phys.*, **98**, 2170 (1993).
42. G. Rauhut and P. Pulay, *LMP2*, University of Arkansas, 1995. LMP2 is part of TEXAS-95, P. Pulay and K. Wolinski, University of Arkansas, 1995.
43. J. W. Boughton and P. Pulay, *J. Comput. Chem.*, **14**, 736 (1993).
44. J. Pipek and P. G. Mezey, *J. Chem. Phys.*, **90**, 4916 (1989).
45. S. F. Boys, In *Quantum Theory of Atoms, Molecules and the Solid State*, P. O. Löwdin, Ed., Academic Press, New York, 1966, p. 253.
46. T. H. Dunning, *J. Chem. Phys.*, **90**, 1007 (1989).
47. P. C. Hariharan and J. A. Pople, *Theor. Chim. Acta*, **28**, 213 (1973).
48. R. Krishnan, J. S. Binkley, R. Seeger, and J. A. Pople, *J. Chem. Phys.*, **72**, 650 (1980).
49. T. Dunning, *J. Chem. Phys.*, **90**, 1007 (1989).
50. P. Pulay and G. Rauhut, Molecular Quantum Mechanics: Methods and Applications [conference contribution], Cambridge, U.K., 1995.
51. E. Buncel and S. Rajagopal, *J. Org. Chem.*, **54**, 798 (1989).
52. P. G. Jösson and A. Kvick, *Acta Crystallogr.*, **B28**, 1827 (1972).
53. G. Rauhut, unpublished data.